arXiv:1404.4123v1 [cs.DS] 16 Apr 2014

Covering problems in edge- and node-weighted graphs

Takuro Fukunaga*

Abstract

This paper discusses the graph covering problem in which a set of edges in an edge- and node-weighted graph is chosen to satisfy some covering constraints while minimizing the sum of the weights. In this problem, because of the large integrality gap of a natural linear programming (LP) relaxation, LP rounding algorithms based on the relaxation yield poor performance. Here we propose a stronger LP relaxation for the graph covering problem. The proposed relaxation is applied to designing primal-dual algorithms for two fundamental graph covering problems: the prize-collecting edge dominating set problem and the multicut problem in trees. Our algorithms are an exact polynomial-time algorithm for the former problem, and a 2-approximation algorithm for the latter problem, respectively. These results match the currently known best results for purely edge-weighted graphs.

1 Introduction

1.1 Motivation

Choosing a set of edges in a graph that optimizes some objective function under constraints on the chosen edges constitutes a typical combinatorial optimization problem and has been investigated in many varieties. For example, the spanning tree problem seeks an acyclic edge set that spans all nodes in a graph, the edge cover problem finds an edge set such that each node is incident to at least one edge in the set, and the shortest path problem selects an edge set that connects two specified nodes. All these problems seek to minimize the sum of the weights assigned to edges.

This paper discusses several graph covering problems. Formally, the *graph covering problem* is defined as follows in this paper. Given a graph $G = (V, E)$ and family $\mathcal{E} \subseteq 2^E$, find a subset F of E that satisfies $F \cap C \neq \emptyset$ for each $C \in \mathcal{E}$, while optimizing some function depending on F. As indicated above, the popular approaches assume an edge weight function $w\colon E \to \mathbb{R}_+$ is given, where $\mathbb{R}_+$ denotes the set of non-negative real numbers, and seeks to minimize $\sum_{e\in F} w(e)$. On the other hand, we aspire to simultaneously minimize edge and node weights. Formally, we let $V(F)$ denote the set of end nodes of edges in F. Given a graph $G = (V, E)$ and weight function $w\colon E \cup V \to \mathbb{R}_+$, we seek a subset F of E that minimizes $\sum_{e\in F} w(e) + \sum_{v\in V(F)} w(v)$ under the constraints on F. Hereafter, we denote $\sum_{e\in F} w(e)$ and $\sum_{v\in V(F)} w(v)$ by $w(F)$ and $w(V(F))$, respectively.

Most previous investigations of the graph covering problem have focused on edge weights. By contrast, node weights have been largely neglected, except in the problems of choosing node sets,

*National Institute of Informatics, 2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo, Japan. JST, ERATO, Kawarabayashi Large Graph Project, Japan. Email: takuro@nii.ac.jp

such as the vertex cover and dominating set problems. To our knowledge, when node weights have been considered in graph covering problems for choosing edge sets, they have been restricted to the Steiner tree problem or its generalizations, possibly because the inclusion of node weights greatly complicates the problem. For example, the Steiner tree problem in edge-weighted graphs can be approximated within a constant factor (the best currently known approximation factor is 1.39 [5, 15]). Conversely, the Steiner tree problem in node-weighted graphs is known to extend the set cover problem (see [20]), indicating that achieving an approximation factor of $o(\log |V|)$ is NP-hard. The literature is reviewed in Section 2. As revealed later, the inclusion of node weights generalizes the set cover problem in numerous fundamental problems.

However, from another perspective, node weights can introduce rich structure into the above problems. In fact, node weights provide useful optimization problems. The objective function counts the weight of a node only once, even if the node is shared by multiple edges. Hence, the objective function defined from node weights includes a certain subadditivity, which cannot be captured by edge weights.

The aim of the present paper is to give algorithms for fundamental graph covering problems in edge- and node-weighted graphs. In solving the problems, we adopt a basic linear programming (LP) technique. Many algorithms for combinatorial optimization problems are typically designed using LP relaxations. However, in problems with node-weighted graphs, the integrality gap of natural relaxations may be excessively large. Therefore, we propose tighter LP relaxations that preclude unnecessary integrality gaps. We then discuss upper bounds on the integrality gap of these relaxations in two fundamental graph covering problems: the edge dominating set (EDS) problem and multicut problem in trees. We prove upper bounds by designing primal-dual algorithms for both problems. The approximation factors of our proposed algorithms match the current best approximations in purely edge-weighted graphs.

1.2 Problem definitions

The EDS problem covers edges by choosing adjacent edges in undirected graphs. For any edge e, let $\delta(e)$ denote the set of edges that share end nodes with e, including e itself. We say that an edge e *dominates* another edge f if $f \in \delta(e)$, and a set F of edges dominates an edge f if F contains an edge that dominates f. Given an undirected graph $G = (V, E)$, a set of edges is called an EDS if it dominates each edge in E. The EDS problem seeks to minimize the weight of the EDS. In other words, the EDS problem is the graph covering problem with $\mathcal{E} = \{\delta(e) \colon e \in E\}$.

In the multicut problem, an instance specifies an undirected graph $G = (V, E)$ and demand pairs $(s_1, t_1), \ldots, (s_k, t_k) \in V \times V$. A *multicut* is an edge set C whose removal from G disconnects the nodes in each demand pair. This problem seeks a multicut of minimum weight. Let $\mathcal{P}_i$ denote the set of paths connecting s_i and t_i. The multicut problem is equivalent to the graph covering problem with $\mathcal{E} = \bigcup_{i=1}^{k} \mathcal{P}_i$.

Our proposed algorithms for solving these problems assume that the given graph G is a tree. In fact, our algorithms are applicable to the prize-collecting versions of these problems, which additionally specifies a penalty function $\pi \colon \mathcal{E} \to \mathbb{R}_+$. In this scenario, an edge set F is a feasible solution even if $F \cap C = \emptyset$ for some $C \in \mathcal{E}$, but imposes a penalty $\pi(C)$. The objective is to minimize the sum of $w(F)$, $w(V(F))$, and the penalty $\sum_{C \in \mathcal{E} : F \cap C = \emptyset} \pi(C)$. The prize-collecting versions of the EDS and multicut problems are referred to as the *prize-collecting EDS problem* and

the *prize-collecting multicut problem*, respectively.

1.3 Our results

Thus far, the EDS problem has been applied only to edge-weighted graphs. The vertex cover problem can be reduced to the EDS problem while preserving the approximation factors [6]. The vertex cover problem is solvable by a 2-approximation algorithm, which is widely regarded as the best possible approximation. Indeed, assuming the unique game conjecture, Khot and Regev [19] proved that the vertex cover problem cannot be approximated within a factor better than 2. Fujito and Nagamochi [11] showed that a 2-approximation algorithm is admitted by the EDS problem, which matches the approximation hardness known for the vertex cover problem. In the Appendix, we show that the EDS problem in bipartite graphs generalizes the set cover problem if assigned node weights and generalizes the non-metric facility location problem if assigned edge and node weights. This implies that including node weights increases difficulty of the problem even in bipartite graphs.

On the other hand, Kamiyama [18] proved that the prize-collecting EDS problem in an edge-weighted graph admits an exact polynomial-time algorithm if the graph is a tree. As one of our main results, we show that this idea is extendible to problems in edge- and node-weighted trees.

Theorem 1. *The prize-collecting EDS problem admits a polynomial-time exact algorithm for edge- and node-weighted trees.*

Theorem 1 will be proven in Section 4. As demonstrated in the Appendix, the prize-collecting EDS problem in general edge- and node-weighted graphs admits an $O(\log |V|)$-approximation, which matches the approximation hardness on the set cover problem and the non-metric facility location problem.

The multicut problem is hard even in edge-weighted graphs; the best reported approximation factor is $O(\log k)$ [13]. The multicut problem is known to be both NP-hard and MAX SNP-hard [10], and admits no constant factor approximation algorithm under the unique game conjecture [7]. However, Garg, Vazirani, and Yannakakis [14] developed a 2-approximation algorithm for the multicut problem with edge-weighted trees. They also mentioned that, although the graphs are restricted to trees, the structure of the problem is sufficiently rich. They showed that the tree multicut problem includes the set cover problem with tree-representable set systems. They also showed that the vertex cover problem in general graphs is simply reducible to the multicut problem in star graphs, while preserving the approximation factor. This implies that the 2-approximation seems to be tight for the multicut problem in trees. As a second main result, we extended this 2-approximation to edge- and node-weighted trees, as stated in the following theorem.

Theorem 2. *The prize-collecting multicut problem admits a 2-approximation algorithm for edge- and node-weighted trees.*

Both algorithms claimed in Theorems 1 and 2 are primal-dual algorithms, that use the LP relaxations we propose. These algorithms fall in the same frameworks as those proposed in [14, 18] for edge-weighted graphs. However, they need several new ideas to achieve the claimed performance because our LP relaxations are much more complicated than those used in [14, 18].

The remainder of this paper is organized as follows. After surveying related work in Section 2, we define our LP relaxation for the prize-collecting graph covering problem in Section 3. Using

this relaxation, we prove Theorems 1 and 2 in Sections 4 and 5, respectively. The paper concludes with Section 6. In the Appendix, we show that the prize-collecting EDS problem in edge- and node-weighted graphs generalizes the set cover problem and facility location problem, and admits an $O(\log |V|)$-approximation algorithm.

2 Related work

As mentioned in Section 1, the graph covering problem in node-weighted graphs has thus far been applied to the Steiner tree problem and its generalizations. Klein and Ravi [20] proposed an $O(\log |V|)$-approximation algorithm for the Steiner tree problem with node weights. Nutov [25, 26] extended this algorithm to the survivable network design problem with higher connectivity requirements. An $O(\log |V|)$-approximation algorithm for the prize-collecting Steiner tree problem with node weights was provided by Moss and Rabani [23]; however, as noted by Könemann, Sadeghian, and Sanità [21], the proof of this algorithm contains a technical error. This error was corrected in [21]. Bateni, Hajiaghayi, and Liaghat [1] proposed an $O(\log |V|)$-approximation algorithm for the prize-collecting Steiner forest problem and applied it to the budgeted Steiner tree problem. Chekuri, Ene, and Vakilian [8] gave an $O(k^2 \log |V|)$-approximation algorithm for the prize-collecting survivable network design problem with edge-connectivity requirements of maximum value k. Later, they improved their approximation factor to $O(k \log |V|)$, and also extended it to node-connectivity requirements (see [30]). Naor, Panigrahi, and Singh [24] established an online algorithm for the Steiner tree problem with node weights which was extended to the Steiner forest problem by Hajiaghayi, Liaghat, and Panigrahi [16]. The survivable network design problem with node weights has also been extended to a problem called the network activation problem [28, 27, 12].

The prize-collecting EDS problem generalizes the $\{0,1\}$-EDS problem, in which given demand edges require being dominated by a solution edge set. The $\{0,1\}$-EDS problem in general edge-weighted graphs admits a 8/3-approximation, which was proven by Berger et al. [2]. This 8/3-approximation was extended to the prize-collecting EDS problem by Parekh [29]. Berger and Parekh [3] designed an exact algorithm for the $\{0,1\}$-EDS problem in edge-weighted trees, but their result contains an error [4]. Since the prize-collecting EDS problem embodies the $\{0,1\}$-EDS problem, the latter problem could be alternatively solved by an algorithm developed for the prize-collecting EDS problem in edge-weighted trees, proposed by Kamiyama [18].

3 LP relaxations

This section discusses LP relaxations for the prize-collecting graph covering problem in edge and node-weighted graphs.

In a natural integer programming (IP) formulation of the graph covering problem, each edge e is associated with a variable $x(e) \in \{0,1\}$, and each node v is associated with a variable $x(v) \in \{0,1\}$. $x(e) = 1$ denotes that e is selected as part of the solution set, while $x(v) = 1$ indicates the selection of an edge incident to v. In the prize-collecting version, each demand set $C \in \mathcal{E}$ is also associated with a variable $z(C) \in \{0,1\}$, where $z(C) = 1$ indicates that the covering constraint corresponding to C is not satisfied. For $F \subseteq E$, we let $\delta_F(v)$ denote the set of edges incident to v in F. The subscript may be removed when $F = E$. An IP of the prize-collecting graph covering problem is

then formulated as follows.

$$
\begin{aligned}
\text{minimize} \quad & \sum_{e\in E} w(e)x(e) + \sum_{v\in V} w(v)x(v) + \sum_{C\in\mathcal{E}} \pi(C)z(C) & \\
\text{subject to} \quad & \sum_{e\in C} x(e) \geq 1 - z(C) & \text{for } C \in \mathcal{E}, \\
& x(v) \geq x(e) & \text{for } v \in V, e \in \delta(v), \\
& x(e) \geq 0 & \text{for } e \in E, \\
& x(v) \geq 0 & \text{for } v \in V, \\
& z(C) \geq 0 & \text{for } C \in \mathcal{E}.
\end{aligned}
$$

In the above formulation, the first constraints specify the covering constraints, while the second constraints indicate that if the solution contains an edge e incident to v, then $x(v) = 1$. In the graph covering problem (without penalties), z is fixed at 0.

To obtain an LP relaxation, we relax the definitions of x and z in the above IP to $x \in \mathbb{R}_+^{E\cup V}$ and $z \in \mathbb{R}_+^C$. However, this relaxation may introduce a large integrality gap into the graph covering problem with node-weighted graphs, as shown in the following example. Suppose that $\mathcal{E}$ comprises a single edge set C, and each edge in C is incident to a node v. Let the weights of all edges and nodes other than v be 0. In this scenario, the optimal value of the graph covering problem is $w(v)$. On the other hand, the LP relaxation admits a feasible solution x such that $x(v) = 1/|C|$ and $x(e) = 1/|C|$ for each edge $e \in C$. The weight of this solution is $w(v)/|C|$, and the integrality gap of the relaxation for this instance is $|C|$.

This phenomenon occurs even in the EDS problem and multicut problem in trees. For instance, consider a star of n leaves in the EDS problem. The weight of all edges and nodes is 0 except the center node v, whose weight is 1. In this instance of the EDS problem, the weight of any EDS is 1. On the other hand, LP relaxation admits a feasible solution x such that $x(e) = 1/n$ for each edge e, and $x(u) = 1/n$ for each node u. Since the weight of this fractional solution is $1/n$, the integrality gap is n.

Let uv denote an edge that joints nodes u and v. To determine the integrality gap in the multicut problem, we consider that each edge uv in the star is subdivided into two edges us and sv. The subdivision imposes a weight of 1 on node s. All edges and remaining nodes (i.e., the center node and all leaves) have weight 0. All pairs of leaves are demand pairs. A path between the center node and a leaf is called a *leg*. In this instance, any multicut must choose at least one edge from each of $n-1$ legs. Hence, the minimum multicut weight is $n-1$. On the other hand, if $x(e) = 1/4$ for every edge e and $x(v') = 1/4$ for every node v', the weight is $n/4$ (such a fractional solution is feasible to the relaxation). Hence, the integrality gap of the relaxation is at least 4. By contrast, Garg, Vazirani, Yannakakis [14] proved that the integrality gap of the relaxation is at most 2 when node weights are not considered.

The above poor examples can be excluded if the second constraints in the relaxation are replaced by $x(v) \geq \sum_{e\in\delta(v)} x(e)$ for $v \in V$. However, the LP obtained by this modification does not relax the graph covering problem if the optimal solutions contain high-degree nodes. Thus, we introduce a new variable $y(C, e)$ for each pair of $C \in \mathcal{E}$ and $e \in C$, and replace the second constraints by $x(v) \geq \sum_{e\in\delta(v)} y(C, e)$, where $v \in V$ and $C \in \mathcal{E}$. $y(C, e) = 1$ indicates that e is chosen to satisfy the covering constraint of C, and $y(C, e) = 0$ implies the opposite. Roughly speaking, $y(C, \cdot)$ represents

a minimal fractional solution for covering a single demand set C. If a single covering constraint is imposed, the degree of each node is at most one in any minimal integral solution. Then the graph covering problem is relaxed by the LP even after modification. Summing up, we formulate our LP relaxation for an instance $I = (G, \mathcal{E}, w, \pi)$ of the prize-collecting graph covering problem as follows.

$$
\begin{aligned}
&P(I) = \\
&\text{minimize} && \sum_{e \in E} w(e)x(e) + \sum_{v \in V} w(v)x(v) + \sum_{C \in \mathcal{E}} \pi(C)z(C) && \\
&\text{subject to} && \sum_{e \in C} y(C,e) \geq 1 - z(C) && \text{for } C \in \mathcal{E}, \\
& && x(v) \geq \sum_{e \in \delta_C(v)} y(C,e) && \text{for } v \in V, C \in \mathcal{E}, \\
& && x(e) \geq y(C,e) && \text{for } C \in \mathcal{E}, e \in C, \\
& && x(e) \geq 0 && \text{for } e \in E, \\
& && x(v) \geq 0 && \text{for } v \in V, \\
& && y(C,e) \geq 0 && \text{for } C \in \mathcal{E}, e \in C, \\
& && z(C) \geq 0 && \text{for } C \in \mathcal{E}.
\end{aligned}
$$

Theorem 3. *Let I be an instance of the prize-collecting graph covering problem in edge- and node-weighted graphs. $P(I)$ is not greater than the optimal value of I.*

Proof. Let F be an optimal solution of I. We define a solution (x, y, z) of $P(I)$ from F. For each $C \in \mathcal{E}$, we set $z(C)$ to 0 if $F \cap C \neq \emptyset$, and 1 otherwise. If $F \cap C \neq \emptyset$, we choose an arbitrary edge $e \in F \cap C$, and let $y(C, e) = 1$. For the remaining edges e', we assign $y(C, e') = 0$. In this way, the values of variables in y are defined for each $C \in \mathcal{E}$. $x(e)$ is set to 1 if $e \in F$, and 0 otherwise. $x(v)$ is set to 1 if F contains an edge incident to v, and 0 otherwise.

For each $C \in \mathcal{E}$ with $z(C) < 1$, exactly one edge e satisfies $y(C, e) = 1$, and this e is included in $F \cap C$. If $y(C, e) = 1$, then $x(e) = 1$, and each end node v of e satisfies $x(v) = 1$. For a pair of $v \in V$ and $C \in \mathcal{E}$, $y(C, e)$ is one for exactly one edge $e \in \delta_C(v)$, and zero for the remaining edges in $\delta_C(v)$. Thus, (x, y, z) is feasible. The objective value of (x, y, z) in $P(I)$ is given by $w(F) + w(V(F)) + \sum_{C \in \mathcal{E}: F \cap C = \emptyset} \pi(C)$, which is the optimal value of I, and the theorem is proven. □

In some graph covering problems, $\mathcal{E}$ is not explicitly given, and $|\mathcal{E}|$ is not bounded by a polynomial on the input size of the problem. In such cases, the above LP may not be solved in polynomial time because it cannot be written compactly. However, in this scenario, we may define a tighter LP than the natural relaxation if we can find $\mathcal{E}_1, \ldots, \mathcal{E}_t \subseteq \mathcal{E}$ such that $\cup_{i=1}^{t} \mathcal{E}_i = \mathcal{E}$, t is bounded by a polynomial of input size, and the degree of each node is small in any minimal edge set covering all demand sets in $\mathcal{E}_i$ for each $i \in \{1, \ldots, t\}$. Applying these conditions, the present author obtained a new approximation algorithm for solving a problem generalizing some prize-collecting graph covering problems [12].

4 Prize-collecting EDS problem in trees

In this section, we prove Theorem 1. We regard the input graph G as a rooted tree, with an arbitrary node r selected as the root. The *depth* of a node v is the number of edges on the path between r and v. When v lies on the path between r and another node u, we say that v is an *ancestor* of u and u is a *descendant* of v. If the depth of node v is the maximum among all ancestors of u, then v is defined as the *parent* of u. If v is the parent of u, then u is a *child* of v. The upper and lower end nodes of an edge e are denoted by u_e and l_e, respectively. We say that an edge e is an ancestor of a node v and v is a descendant of e when $l_e = v$ or l_e is an ancestor of v. Similarly, an edge e is a descendant of a node v and v is an ancestor of e if $v = u_e$ or v is an ancestor of u_e. An edge e is defined as an ancestor of another edge f if e is an ancestor of u_f.

Recall that $\mathcal{E} = \{\delta(e) \colon e \in E\}$ in the EDS problem. Let $I = (G, w, \pi)$ be an instance of the prize-collecting EDS problem. We denote $\bigcup_{e \in \delta(v)} \delta(e)$ by $\delta'(v)$ for each $v \in V$. Then the dual of $P(I)$ is formulated as follows.

$$
\begin{aligned}
&D(I) = \\
&\text{maximize} && \sum_{e \in E} \xi(e) && && \\
&\text{subject to} && \sum_{e \in \delta(e')} \nu(e', e) \le w(e') && \text{for } e' \in E, && (1)\\
& && \sum_{e \in \delta'(v)} \mu(v, e) \le w(v) && \text{for } v \in V, && (2)\\
& && \xi(e) \le \mu(u, e) + \mu(v, e) + \nu(e', e) && \text{for } e \in E, e' = uv \in \delta(e), && (3)\\
& && \xi(e) \le \pi(e) && \text{for } e \in E, && (4)\\
& && \xi(e) \ge 0 && \text{for } e \in E, && \\
& && \nu(e', e) \ge 0 && \text{for } e' \in E, e \in \delta(e'), && \\
& && \mu(v, e) \ge 0 && \text{for } v \in V, e \in \delta'(v). &&
\end{aligned}
$$

For an edge set $F \subseteq E$, let $\tilde{F}$ denote $\{e \in E \colon \delta_F(e) = \emptyset\}$, and let $\pi(\tilde{F})$ denote $\sum_{e \in \tilde{F}} \pi(e)$. For the instance I, our algorithm yields a solution $F \subseteq E$ and a feasible solution (ξ, ν, μ) to $D(I)$, both satisfying

$$w(F) + w(V(F)) + \pi(\tilde{F}) \le \sum_{e \in E} \xi(e). \tag{5}$$

Since the right-hand side of (5) is at most $P(I)$, F is an optimal solution of I. We note that the dual solution (ξ, ν, μ) is required only for proving the optimality of the solution and need not be computed.

The algorithm operates by induction on the number of nodes of depth exceeding one. In the base case, all nodes are of depth one, indicating that G is a star centered at r. The alternative case is divided into two sub-cases: Case A, in which a leaf edge e of maximum depth satisfies $\pi(e) > 0$; and Case B, which contains no such leaf edge.

Base case

In the base case, G is a star centered at r. Note that all edges in this graph are adjacent. Let $\alpha_1 = \min_{rv \in E}\{w(rv)+w(r)+w(v)\}$ and $\alpha_2 = \sum_{e \in E} \pi(e)$. An edge rv attaining $\alpha_1 = w(rv)+w(r)+w(v)$ is denoted by $e^* = rv^*$.

If $\alpha_1 \geq \alpha_2$, our algorithm sets F as $\emptyset$, and defines $\xi(e)$ as $\pi(e)$ for each $e \in E$. Otherwise (i.e., $\alpha_1 < \alpha_2$), it specifies F as $\{e^*\}$, and sets ξ so that $\sum_{e \in E} \xi(e) = \alpha_1$, and $\xi(e) \leq \pi(e)$ for each $e \in E$, which is possible because $\sum_{e \in E} \pi(e) = \alpha_2 > \alpha_1$. Note that F and ξ defined in this way satisfy (5).

To completely define the dual solution, we must define variables ν and μ. Let $rv \in E$. To satisfy (3) for $e \in E$ and rv, the sum of $\mu(r, e)$, $\mu(v, e)$, and $\nu(rv, e)$ cannot be smaller than $\xi(e)$ for each $e \in E$. Note that in (1), $\nu(rv, e)$ is bounded from above for rv, while in (2) for v, $\mu(r, e)$ and $\mu(v, e)$ are bounded for r and v, respectively. As an alternative interpretation, each $rv \in E$ has capacity $w(rv)$ shared by $\nu(rv, e)$, $e \in \delta(rv)$, and each $v \in V$ has capacity $w(v)$ shared by $\nu(v, e)$, $e \in \delta'(v)$. The following lemma claims that ν and μ may be set to satisfy all of these constraints.

Lemma 1. *Suppose that G is a star. If $\alpha_1 \geq \alpha_2$, let $\xi(e) = \pi(e)$ for each $e \in E$. Otherwise, suppose that ξ is defined to satisfy $\sum_{e \in E} \xi(e) = \alpha_1$, and $\xi(e) \leq \pi(e)$ for each $e \in E$. Then there exists a feasible solution (ξ, ν, μ) to $D(I)$.*

Proof. First, we appropriately define ν and μ. All variables of ν and μ are initialized to 0. We fix an arbitrary ordering of edges in E, and denote the i-th edge by e_i.

Let $rv \in E$. We sequentially select edges e_1 to $e_{|E|}$. On selection of e_i, we first increase $\mu(r, e_i)$ until the increase reaches $\xi(e_i)$ or (2) becomes tight for r. If (2) is tightened for r before $\mu(r, e_i)$ is increased by $\xi(e_i)$, then $\nu(rv, e_i)$ is increased until the total increase reaches $\xi(e_i)$ or (1) becomes tight for rv. Once (1) has tightened for rv, $\mu(v, e_i)$ is increased. The current iteration is terminated when the total increase reaches $\xi(e_i)$. If $i < |E|$, the algorithm advances to the next iteration, and processes e_{i+1}. Since $\sum_{i=1}^{|E|} \xi(e_i) \leq \min\{\alpha_1, \alpha_2\} \leq w(rv) + w(r) + w(v)$, all edges in E can be processed before (2) becomes tight for v.

The above process defines $\nu(rv, e_i)$, $\mu(v, e_i)$, and $\mu(r, e_i)$ for each $i \in \{1, \ldots, |E|\}$. This process is repeated for all $rv \in E$, but $\mu(r, e_i)$ is not increased beyond the first iteration. Note that $\mu(r, e_i)$ is assigned the same value regardless of which edge rv we begin with. Thus, ν and μ have been completely defined, and the feasibility of (ξ, ν, μ) follows from their definitions. □

Case A

In this case, a leaf edge e of maximum depth satisfies $\pi(e) > 0$. Since Case A is not the base case, the depth of l_e exceeds one. Let u denote the upper end node of e. Also let v_0 denote the parent of u, and let $v_1, \ldots, v_k$ be the children of u. Throughout this paper, the sets $\{1, \ldots, k\}$ and $\{0, \ldots, k\}$ are denoted by $[k]$, and $[k]^*$, respectively. The edge joining v_i and u is called e_i, with $i \in [k]^*$ (e is included in $\{e_1, \ldots, e_k\}$). The relationships between these nodes and edges are illustrated in Figure 1. We define $\beta_1 = \min_{i=0}^{k}(w(e_i) + w(u) + w(v_i))$, $\beta_2 = \sum_{i=1}^{k} \pi(e_i)$, and $\beta = \min\{\beta_1, \beta_2\}$. Let $i^* \in [k]^*$ be the index of an edge e_{i^*} that attains $\beta_1 = w(e_{i^*}) + w(u) + w(v_{i^*})$.

The algorithm constructs an instance $I' = (G', w', \pi')$ as follows. Suppose that $\beta_1 > \beta_2$. In this

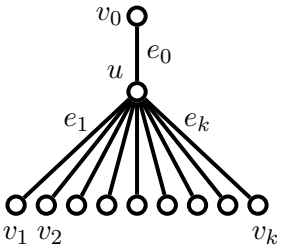

Figure 1: Edges and nodes in Case A

case, G' is defined as G, and $\pi'\colon E' \to \mathbb{R}_+$ is defined by

$$\pi'(e) = \begin{cases} 0 & \text{if } e \in \{e_1, \ldots, e_k\}, \\ \pi(e) & \text{otherwise.} \end{cases}$$

For $\psi \in \mathbb{R}$, we denote $\max\{0, \psi\}$ by $(\psi)_+$. The weight function $w'\colon V' \cup E' \to \mathbb{R}_+$ is defined by

$$w'(e) = \begin{cases} (w(e_i) - (\beta - w(u))_+)_+ & \text{if } e = e_i, i \in [k]^*, \\ w(e) & \text{otherwise} \end{cases}$$

for each $e \in E'$, and

$$w'(v) = \begin{cases} (w(u) - \beta)_+ & \text{if } v = u, \\ w(v_i) - (\beta - w(u) - w(e_i))_+ & \text{if } v = v_i, i \in [k]^*, \\ w(v) & \text{otherwise} \end{cases}$$

for each $v \in V'$. If $\beta_1 \leq \beta_2$, then $G' = (V', E')$ is defined as the tree obtained by removing nodes $v_1, \ldots, v_k$ and edges $e_1, \ldots, e_k$ from G, and $\pi'\colon E' \to \mathbb{R}_+$ is defined by

$$\pi'(e) = \begin{cases} 0 & \text{if } e = e_0, \\ \pi(e) & \text{otherwise.} \end{cases}$$

w' is defined identically to the case $\beta_1 > \beta_2$, ignoring $w'(v_i)$ and $w'(e_i)$ for $i \in [k]$.

If $\beta_1 \leq \beta_2$, the number of nodes with depth exceeding one is lower in G' than in G. Hence, the algorithm inductively finds a solution F' to I' and a feasible dual solution (ξ', ν', μ') to $D(I')$ that satisfy (5). Otherwise, the number of leaf edges e of maximum depth with $\pi'(e) > 0$ is lower in G' than in G. If G' lacks edges of this type, then instance I' is categorized into Case B, and F' and (ξ', ν', μ') are found as demonstrated below. If such edges do exist in G', the algorithm finds F' and (ξ', ν', μ') by induction on the number of such edges. Therefore, it suffices to show that the required F and (ξ, ν, μ) can be constructed from F' and (ξ', ν', μ'), provided that F' and (ξ', ν', μ') exist.

We now define F and (ξ, ν, μ). F is defined by

$$F = \begin{cases} F' \cup \{e_0\} & \text{if } \delta_{F'}(v_0) \neq \emptyset, \beta > w(u) + w(e_0), \\ F' & \text{if } \delta_{F'}(v_0) = \emptyset \text{ or } \beta \leq w(u) + w(e_0), \beta_1 > \beta_2, \\ F' \cup \{e_{i^*}\} & \text{if } \delta_{F'}(v_0) = \emptyset \text{ or } \beta \leq w(u) + w(e_0), \beta_1 \leq \beta_2. \end{cases}$$

Lemma 2. *There exists a feasible solution (ξ, ν, μ) to $D(I)$ that satisfies (5) with F.*

Proof. We first consider the case of $\beta_1 > \beta_2$. In this case, $\xi'(e_i) = 0$ follows from $\pi'(e_i) = 0$ for each $i \in [k]$. We define $\xi(e_i)$ as $\pi(e_i)$ for $i \in [k]$. We also define $\nu(e_j, e_i)$, $\mu(v_j, e_i)$, and $\mu(u, e_i)$ such that $\nu(e_j, e_i) + \mu(v_j, e_i) + \mu(u, e_i) = \xi(e_i)$ holds for each $j \in [k]^*$ and $i \in [k]$. The other dual variables are set to their values assigned in (ξ', ν', μ'). Note that, for each $i \in [k]^*$, we have $w(u) + w(e_i) + w(v_i) - w'(u) - w'(e_i) - w'(v_i) = \beta_2 = \sum_{i=1}^{k} \xi(e_i)$. Hence, $\nu(e_j, e_i)$, $\mu(v_j, e_i)$, and $\mu(u, e_i)$ can be defined without violating (1) or (2) as follows. We sequentially collect edges e_1 to e_k. On selection of e_i, $\mu(u, e_i)$ is increased until the total increase reaches $\xi(e_i)$ or (2) becomes tight for u. If (2) is tightened for u before $\mu(u, e_i)$ has increased by $\xi(e_i)$, then $\nu(e_j, e_i)$ is simultaneously increased for all $j \in [k]^*$ until (1) becomes tight for e_j. Once (1) has tightened for e_j, $\mu(v_j, e_i)$ is increased instead of $\nu(e_j, e_i)$.

We next consider the case of $\beta_1 \leq \beta_2$. In this scenario, $\xi'(e_0) = 0$ holds because $\pi'(e_0) = 0$. We define $\xi(e_i)$, $i \in [k]$ such that $\xi(e_i) \leq \pi(e_i)$ for each $i \in [k]$ and $\sum_{i=1}^{k} \xi(e_i) = \beta_1$, which is possible because $\sum_{i=1}^{k} \pi(e_i) = \beta_2 \geq \beta_1$. $\nu(e_i, e_0)$ and $\mu(v_i, e_0)$ are set to 0 for each $i \in [k]$. We also define $\nu(e_j, e_i)$, $\mu(v_j, e_i)$, and $\mu(u, e_i)$ such that $\nu(e_j, e_i) + \mu(v_j, e_i) + \mu(u, e_i) = \xi(e_i)$ for each $j \in [k]^*$ and $i \in [k]$ as specified for $\beta_1 > \beta_2$. The other variables are set to their values assigned in (ξ', ν', μ'). The feasibility of (ξ, ν, μ) follows from its definition.

We now prove that F and ξ satisfy (5). Without loss of generality, we can assume $|\delta_{F'}(u)| \leq 1$ (if this condition is false, we can remove edges e_i, where $i \in [k]$, from F' until $|\delta_{F'}(u)| = 1$). The objective value of F exceeds that of F' by at most β, unless $e_0 \in F \setminus F'$ and $i^* \neq 0$. If $e_0 \in F \setminus F'$ and $i^* \neq 0$, then $\delta_{F'}(v_0) \neq \emptyset$ and $\beta > w(u) + w(e_0)$ by the definition of F. Since $\delta_{F'}(v_0) \neq \emptyset$, $w'(v_0)$ is counted in the objective value of F'. Thus, the objective values increases from F' to F by $w(v_0) - w'(v_0) + w(e_0) + w(u)$. From $\beta > w(u) + w(e_0)$, it follows that $w'(u) = w'(e_0) = 0$, therefore, the objective function increases by β. Since $\sum_{e \in E} \xi(e) - \sum_{e \in E'} \xi(e) = \beta$, (5) is satisfied. □

Case B

In this case, $\pi(e) = 0$ holds for all leaf edges e of maximum depth. Let s be the grandparent of a leaf node of maximum depth. Also, let $u_1, \ldots, u_k$ be the children of s, and e_i be the edge joining s and u_i for $i \in [k]$. In the following discussion, we assume that s has a parent, and that each node u_i has at least one child. This discussion is easily modified to cases in which s has no parent or some node u_i has no child. We denote the parent of s by u_0, and the edge between u_0 and s by e_0. For each $i \in [k]$, let V_i be the set of children of u_i, and H_i be the set of edges joining u_i to its child nodes in V_i. Also define $h_i = u_i v_i$ as an edge that attains $\min_{u_i v \in H_i}(w(u_i v) + w(v))$. The relationships between these nodes and edges are illustrated in Figure 2.

Now define $\theta_1 = \min_{i=0}^{k}(w(e_i) + w(u_i) + w(s))$, $\theta_2 = \sum_{i=1}^{k} \min\{w(u_i) + w(v_i) + w(h_i), \pi(e_i)\}$, and let $\theta = \min\{\theta_1, \theta_2\}$. We denote the index $i \in [k]$ of an edge e_i that attains $\theta_1 = w(e_i) + w(u_i) + w(s)$ by i^*, and specify $K = \{i \in [k] \colon w(u_i) + w(v_i) + w(h_i) \leq \pi(e_i)\}$.

We define $I' = (G', w', \pi')$ as follows. If $\theta_1 \geq \theta_2$, then G' is the tree obtained by removing all edges in $\bigcup_{i \in [k]} H_i$ and all nodes in $\bigcup_{i \in [k]} V_i$ from G, and $\pi' \colon E' \to \mathbb{R}_+$ is defined such that

$$\pi'(e) = \begin{cases} 0 & \text{if } e \in \{e_1, \ldots, e_k\}, \\ \pi(e) & \text{otherwise} \end{cases}$$

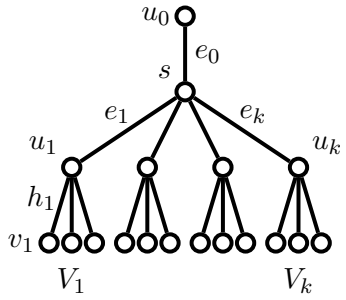

Figure 2: Edges and nodes in Case B

for $e \in E'$. In this case, $w' \colon V' \cup E' \to \mathbb{R}_+$ is defined by

$$w'(v) = \begin{cases} (w(s) - \theta)_+ & \text{if } v = s, \\ w(u_i) - (\theta - w(s) - w(e_i))_+ & \text{if } v = u_i, i \in [k]^* \\ w(v) & \text{otherwise} \end{cases}$$

for $v \in V'$, and

$$w'(e) = \begin{cases} (w(e_i) - (\theta - w(s))_+)_+ & \text{if } e = e_i, i \in [k]^*, \\ w(e) & \text{otherwise,} \end{cases}$$

for $e \in E'$. If $\theta_1 < \theta_2$, then $e_1, \ldots, e_k$, and their descendants are removed from G to obtain G', and π' is defined by

$$\pi'(e) = \begin{cases} 0 & \text{if } e = e_0, \\ \pi(e) & \text{otherwise.} \end{cases}$$

Moreover, w' for E' and V' is defined as in the case $\theta_1 \geq \theta_2$, disregarding the weights of edges and nodes removed from G'.

Since G' has fewer nodes of depth exceeding one than G, the algorithm inductively finds a solution F' to I', and a feasible solution (ξ', ν', μ') to $D(I')$ satisfying (5). F is constructed from F' as follows.

$$F = \begin{cases} F' \cup \{e_0\} & \text{if } \delta_{F'}(u_0) \neq \emptyset, \theta > w(s) + w(e_0), \\ F' & \text{if } \delta_{F'}(u_0) = \emptyset \text{ or } \theta \leq w(s) + w(e_0), \delta_{F'}(s) \neq \emptyset, \\ F' \cup \{h_i \colon i \in K\} & \text{if } \delta_{F'}(u_0) = \emptyset \text{ or } \theta \leq w(s) + w(e_0), \delta_{F'}(s) = \emptyset, \theta_1 \geq \theta_2, \\ F' \cup \{e_{i^*}\} & \text{if } \delta_{F'}(u_0) = \emptyset \text{ or } \theta \leq w(s) + w(e_0), \delta_{F'}(s) = \emptyset, \theta_1 < \theta_2. \end{cases}$$

We define $\xi(e_1), \ldots, \xi(e_k)$ such that $\xi(e_i) \leq \min\{w(u_i) + w(v_i) + w(h_i), \pi(e_i)\}$ for $i \in [k]$ and $\sum_{i=1}^k \xi(e_i) = \theta$, which is possible because $\sum_{i=1}^k \min\{w(u_i) + w(v_i) + w(h_i), \pi(e_i)\} = \theta_2 \geq \theta$. We also define $\xi(e) = 0$ for each $e \in \bigcup_{i=1}^k H_i$. The other variables in ξ are set to their values in ξ'. The following lemma states that this ξ can form a feasible solution to $D(I)$.

Lemma 3. *Suppose that $\xi(e_1), \ldots, \xi(e_k)$ satisfy $\xi(e_i) \leq \min\{w(u_i) + w(v_i) + w(h_i), \pi(e_i)\}$ for each $i \in [k]$ and $\sum_{i=1}^k \xi(e_i) = \theta$. Further, suppose that $\xi(e) = 0$ holds for each $e \in \bigcup_{i=1}^k H_i$, and the other variables in ξ are set to their values in ξ'. Then there exist ν and μ such that (ξ, ν, μ) is feasible to $D(I)$.*

Proof. For $i \in [k]$ and $v \in V_i$, we define $\mu(v, e_i)$ and $\nu(u_i v, e_i)$ such that $\mu(v, e_i) + \nu(u_i v, e_i) = \min\{w(v_i) + w(h_i), \xi(e_i)\}$. This may be achieved without violating the constraints, because $w(v) + w(u_i v) \geq w(v_i) + w(h_i)$. We also define $\nu(u_i v, e_i)$ as $(\xi(e_i) - w(u_i) - w(h_i))_+$. These variables satisfy (1) for $u_i v$, (2) for v and u_i, and (3) for $(e_i, u_i v)$. $\nu(e_j, e_i)$ for $i \in [k]$ and $j \in [k]^*$, and $\mu(v, e_i)$ for $i \in [k]$ and $v \in \{s\} \cup \{u_j \colon j \in [k]^*, j \neq i\}$ are set to 0. The other variables in ν and μ are set to their values in ν' and μ'. To advance the proof, we introduce an algorithm that increases $\nu(e_j, e_i)$ for $i \in [k]$ and $j \in [k]^*$, and $\mu(v, e_i)$ for $i \in [k]$ and $v \in \{s, u_0, \ldots, u_k\}$. At the completion of the algorithm, (ξ, ν, μ) is a feasible solution to $D(I)$.

The algorithm performs k iterations, and the i-th iteration increases the variables to satisfy (3) for each pair of e_i and e_j, where $j \in [k]^*$. The algorithm retains a set `Var` of variables to be increased. We introduce a notion of time: Over one unit of time, the algorithm simultaneously increases all variables in `Var` by one. The time consumed by the i-th iteration is $\xi(e_i)$.

At the beginning of the i-th iteration, `Var` is initialized to $\{\mu(u_j, e_i) \colon j \in [k]^*\}$. The algorithm updates `Var` during the i-th iteration as follows.

- At time $(\xi(e_i) - w(v_i) - w(h_i))_+$, $\mu(u_i, e_i)$ is added to `Var` if `Var` $\neq \{\mu(s, e_i)\}$;
- If (2) becomes tight for u_j under the increase of $\mu(u_j, e_i) \in$ `Var`, then $\mu(u_j, e_i)$ is replaced by $\nu(e_j, e_i)$ for each $j \in [k]^*$;
- If (1) becomes tight for e_j under the increase of $\nu(e_j, e_i) \in$ `Var` with some $j \in [k]^*$, then `Var` is reset to $\{\mu(s, e_i)\}$.

We note that the time spent between two consecutive updates may be zero.

`Var` always contains a variable that appears in the right-hand side of (3) for (e_i, e_j) with $j \in [k]^* \setminus \{i\}$, and for (e_i, e_i) after time $(\xi(e_i) - w(v_i) - w(h_i))_+$. The algorithm updates `Var` so that (1) and (2) hold for all variables except s. Hence, to show that (ξ, ν, μ) is a feasible solution to $D(I)$, it suffices to show that (2) for s does not become tight before the algorithm is completed.

We complete the proof by contradiction. Suppose that (2) for s tightens at time $\tau < \xi(e_i)$ in the i-th iteration. Since `Var` $= \{\mu(s, e_i)\}$ at this moment, there exists $j \in [k]^*$ such that (1) for e_j and (2) for u_j are tight. The variables in the left-hand sides of (1) for e_j and (2) for u_j and s are not simultaneously increased. Nor are these variables increased over time $(\xi(e_j) - w(v_j) - w(h_j))_+$ in the j-th iteration, and $\mu(u_j, e_j)$ is initialized to $(\xi(e_j) - w(v_j) - w(h_j))_+$. From this argument, it follows that $w(s) + w(u_j) + w(e_j) < \sum_{i'=1}^{k} \xi(e_{i'}) \leq \theta$. However, this result is contradicted by the definition of θ, which implies that $\theta \leq \theta_1 \leq w(s) + w(u_j) + w(e_j)$. Thus, the claim is proven. □

Lemma 4. *F and ξ satisfy* (5).

Proof. For each $i \in [k]$, either $e_i \notin E'$ holds, or $\xi'(e_i) = 0$ holds (because $\pi'(e_i) = 0$). Hence, $\sum_{e \in E} \xi(e) = \sum_{i=1}^{k} \xi(e_i) + \sum_{e \in E'} \xi'(e) = \theta + \sum_{e \in E'} \xi'(e)$. Therefore, it suffices to prove that $\sum_{e \in F} w(e) \leq \theta + \sum_{e \in F'} w'(e)$.

Without loss of generality, we can assume $|\delta_{F'}(e_0)| \leq 1$ (if false, we can remove edges e_i, $i \in [k]^*$ from F' until $|\delta_{F'}(e_0)| = 1$). In the sequel, we discuss only the case of $\delta_{F'}(u_0) \neq \emptyset$ and $\theta > w(s) + w(e_0)$. In the alternative case, the claim immediately follows from the definitions of F and w'. $\delta_{F'}(u_0) \neq \emptyset$ implies that $w'(u_0)$ is counted in the objective value of F'. Moreover, $w'(s) = w'(e_0) = 0$ follows from $\theta > w(s) + w(e_0)$. Thus, the objective values increase from F' to F by $w(u_0) - w'(u_0) + w(e_0) + w(s)$, which equals θ. □

Proof of Theorem 1. We have proven that our algorithm always finds a solution F to I and a feasible solution (ξ, ν, μ) to $D(I)$, where both solutions satisfy (5). By the duality of LPs, the right-hand side of (5) cannot exceed the optimal value of I; thus, F is an optimal solution to I. □

5 Multicut problem in trees

In this section, we prove Theorem 2. Again the input tree G is rooted by selecting an arbitrary root node. For each $i \in [k]$, we let P_i denote the path connecting s_i and t_i, and lca_i denote the maximum-depth common ancestor of s_i and t_i. The paths $P_1, \ldots, P_k$ are called *demand paths*. We also denote the set of edges in P_i by E_i, and the set of nodes in P_i by V_i for notational convenience. We say that an edge e *covers* a demand path P_i if $e \in E_i$. The multicut problem in G seeks a minimum weight set of edges that covers all demand paths.

The prize-collecting multicut problem can be reduced to the multicut problem as follows. For each $i \in [k]$, add new nodes s'_i, s''_i and new edges $s_i s'_i, s'_i s''_i$ to G, and replace the i-th demand pair by (s''_i, t_i). Those new nodes and edges are weighted by $w(s'_i s''_i) = \pi(i)$, $w(s_i s'_i) = +\infty$, and $w(s'_i) = w(s''_i) = 0$. Choosing $s'_i s''_i$ into a solution to this new instance of the multicut problem corresponds to violating the i-th demand in the original instance of the prize-collecting multicut problem.

Due to this reduction, we consider only the multicut problem in trees, which is equivalent to assuming that $\pi(i) = +\infty$ for all $i \in [k]$. For an instance $I = (G, w)$ of the multicut problem, the dual of the LP relaxation $P(I)$ is given by

$$
\begin{aligned}
&D(I) = \\
&\text{maximize} && \sum_{i \in [k]} \xi(i) && \\
&\text{subject to} && \xi(i) \le \nu(e,i) + \mu(u_e, i) + \mu(l_e, i) && \text{for } i \in [k], e \in E_i, && (6)\\
& && \sum_{i \in [k] : e \in E_i} \nu(e,i) \le w(e) && \text{for } e \in E, && (7)\\
& && \sum_{i \in [k] : v \in V_i} \mu(v,i) \le w(v) && \text{for } v \in V, && (8)\\
& && \xi(i) \ge 0 && \text{for } i \in [k], && \\
& && \nu(e,i) \ge 0 && \text{for } i \in [k], e \in E_i, && \\
& && \mu(v,i) \ge 0 && \text{for } i \in [k], v \in V_i. &&
\end{aligned}
$$

Our algorithm initializes the solution set F to an empty set, and the dual solution (ξ, ν, μ) to 0. The algorithm proceeds in two phases; the increase phase and deletion phase. The algorithm iterates in the increase phase, selecting edges covering demand paths not previously covered by F and adding them to F while updating the dual solution. The increase phase terminates when all demand paths have been covered by F. In the deletion phase, F is converted into a minimal solution by removing some edges.

The demand pairs are assumed to be sorted in the decreasing order of depth of lca_i, implying that lca_i is not a descendant of lca_j if $j < i$.

Increase phase

At the beginning of each iteration in the increase phase, the algorithm selects the minimum index i for which P_i is not covered by the current solution F. It then updates the dual solution (ξ, ν, μ), and adds several edges to F, one of which covers P_i. If all demand paths are covered by F after this operation, the increase phase is terminated and the algorithm proceeds to the deletion phase; otherwise, it begins the next iteration. The iteration that processes P_i is called *the iteration for* i. In the following discussion, we explain the update process of the dual solution, and how edges are selected for addition to F in the iteration for i.

First, we define some terminologies. We say that $e \in E_i$ is *tight* with regard to $i \in [k]$ if (6) becomes an equality for (i, e). We say that e is a *bottleneck edge* with regard to i if it is tight with regard to i, if (7) becomes an equality for e, and if (8) becomes an equality for both end nodes of e.

At the beginning of the iteration, ν is assumed to be minimal under the condition that (ξ, ν, μ) is a feasible solution to $D(I)$. This condition can be assumed without loss of generality because arbitrarily decreasing ν makes it minimal. By this assumption, if $\nu(e, j) > 0$ for some $j \in [i-1]$ and $e \in E_j$, then e is tight with regard to j.

The algorithm attempts to continuously increase $\xi(i)$. As in Section 4, we introduce a time interval, during which $\xi(i)$ increases by one. To satisfy (6), $\nu(e, i)$, $\mu(u_e, i)$ or $\mu(l_e, i)$ must be increased at the speed of $\xi(i)$ for each edge $e \in E_i$ that is tight with regard to i. If no bottleneck edge exists with regard to i, the algorithm retains an edge set $H \subseteq \{e \in E_i \colon \sum_{j \in [k]: e \in E_j} \nu(e, j) < w(e)\}$ and a node set $U \subseteq \{v \in V_i \colon \sum_{j \in [k]: v \in V_j} \mu(v, j) < w(v)\}$ such that each tight edge in E_i is included in H or is incident to a node in U. The algorithm increases $\nu(e, i)$, $e \in H$, and $\mu(v, i)$, $v \in U$ at the same speed as $\xi(i)$. We note that H and U are computed greedily so that they are minimal.

We now explain how the algorithm handles a bottleneck edge e. For an end node v of the bottleneck edge e, we define $J(i, v)$ as $\{j \in [i-1] \colon \mu(v, j) > 0\}$. The algorithm attempts to decrease $\mu(v, j)$, defined at an end node v of e and $j \in J(i, v)$. Below we detail how $\mu(v, j)$ is decreased while retaining the feasibility of (ξ, μ, ν). We note that decrease of $\mu(v, j)$ is not always possible. We call v *relaxable* (with regard to i) if $\mu(v, j)$ can be decreased for some $j \in J(i, v)$. If E_i contains bottleneck edges, the algorithm maintains

- a set R of relaxable nodes such that each bottleneck edge $e \in E_i$ is incident to at least one node in R,
- an edge set $H \subseteq \{e \in E_i \colon \sum_{j \in [k]: e \in E_j} \nu(e, j) < w(e)\}$,
- and a node set $U \subseteq \{v \in V_i \colon \sum_{j \in [k]: v \in V_j} \mu(v, j) < w(v)\}$.

R, H, and U are minimal under the condition that each tight edge is included in H, or is incident to a node in $R \cup U$. The algorithm increases $\xi(i)$, $\nu(e, i)$ for $e \in H$, and $\mu(v, i)$ for $v \in U \cup R$ at the same speed, where increasing $\mu(v, i)$ for $v \in R$ involves decreasing $\mu(v, j)$ for some $j \in J(i, v)$ and updating other variables, as explained below.

We now explain how $\mu(v, j)$ is decreased for some $j \in J(i, v)$, and formally define the relaxability of v. $\mu(v, j) > 0$ implies that E_j contains one or two edges incident to v. Suppose that E_j contains a single edge, f. Let u be the other end node of f. If f is not tight with regard to j, then $\mu(v, j)$ is decreased until f becomes tight. Even if f is tight with regard to j, $\nu(f, j)$ or $\mu(u, j)$ is

increased while $\mu(v, j)$ is decreased at the same speed, provided that f is not a bottleneck edge with regard to j. This action retains the feasibility because (7) for f or (8) for u is not tight unless f is a bottleneck edge with regard to j. If f is a bottleneck edge with regard to j, the algorithm recursively attempts to decrease $\mu(u, j')$ for some $j' \in J(j, u)$ if u is relaxable, and increase $\mu(u, j)$. Under these update rules, $\mu(v, j)$ is decreased without violating the feasibility of (ξ, μ, ν). If E_j contains two edges f and f' incident to v, $\mu(v, j)$ decreases only when allowed for both f and f'. We define v as relaxable if one of these updates is possible. v is not relaxable under the following conditions.

Fact 1. *A node $v \in V_i$ is not relaxable with regard to i if and only if $J(i, v) = \emptyset$, or for each $j \in J(i, v)$, E_j contains a bottleneck edge f whose other end node is non-relaxable and which is incident to v.*

We note that decreasing $\mu(v, j)$ for some relaxable node $v \in V_i$ and $j \in J(i, v)$ may cause other variables to increase. In this case, if P_j shares nodes or edges with P_i, increasing $\xi(i)$ by $\epsilon > 0$ may increase the left-hand sides of (7) and (8) by more than ϵ. Hence, ϵ must be set sufficiently small that the feasibility of (ξ, μ, ν) is maintained. In implementing the increase phase, we recommend solving an LP for deciding the increment of $\xi(i)$ in a single step. The maximum increment ϵ for $\xi(i)$ can be computed by formulating the problem as an LP.

When $\xi(i)$ increases no further, the algorithm adds several edges to F. At this moment, E_i includes a bottleneck edge e such that e is tight with regard to all $j \in [k]$, and neither of its end nodes are relaxable. If two or more such edges exist, the edge of maximum depth, denoted e, is added to F. We call e the *witness* of P_i.

The algorithm then completes the following operations for each end node v of e. By Fact 1, E_j contains a bottleneck edge f incident to v for each $j \in J(i, v)$, where $f = e$ possibly holds. The algorithm adds such f to F for each $j \in J(i, v)$ with $e \notin E_j$. Since the other end node v' of f is non-relaxable, $E_{j'}$ also contains a bottleneck edge f' incident to v' for each $j' \in J(j, v')$. If f is added to F and $f \notin E_{j'}$, the algorithm adds each of such f' to F and repeats the process for the other end nodes of f'.

Lemma 5. *Let $e \in F$. At the completion of the increase phase, e is a bottleneck edge with regard to each $i \in [k]$. Moreover, neither end node of e is relaxable with regard to each $i \in [k]$.*

Proof. When e is added to F, it satisfies the above conditions. In later iterations, the algorithm does not decrease $\nu(e', i)$ for any bottleneck edge e' nor $\mu(v, i)$ for an end node v of e'. Hence, e satisfies the conditions at completion of the increase phase. □

Deletion phase

Let I be the set of indices i for which P_i was considered in the increase phase. In other words, i was the minimum index for which P_i was covered by no edge in F at the beginning of some iteration of the increase phase. The deletion phase is also iterated, sequentially processing P_i in the decreasing order of $i \in I$. As defined in the increase phase, the iteration that considers P_i is called *the iteration for i*. Briefly, the deletion phase selects edges from F to obtain a final solution F' in which each $\xi(i)$, $i \in I$ contributes to at most two edges.

F' is initialized to the empty set. Suppose that the algorithm iterates $i \in I$. We denote $\{v \in V_i \colon \mu(v,i) > 0, \delta_{F'}(v) \neq \emptyset\}$ by $\tilde{V}_i$. Let v be a node in $\tilde{V}_i$ with no ancestor in $\tilde{V}_i$. By Lemma 5, $\delta_F(v) \supseteq \delta_{F'}(v) \neq \emptyset$ implies that v is non-relaxable with regard to i. Suppose now that an edge in $\delta_{F'}(v)$ is added to F' in the iteration for i' with $i' > i$. $\mu(v,i) > 0$ implies that $i \in J(i', v)$, and hence by Fact 1, E_i contains a bottleneck edge f incident to v, whose other end node is non-relaxable. If E_i contains two such edges, and if v is the lower end node of one of those edges, f becomes the edge satisfying $l_f = v$. Otherwise, f is arbitrarily selected from the candidate edges. If f has not previously belonged to F', it is added to F'. If F' contains an edge $f' \in E_i$ which is a descendant of f, f' is deleted from F'.

The above v can be selected from at most two choices in $\tilde{V}_i$. The algorithm completes the above operation for each of these nodes. At the end of the operation, if $F' \cap E_i$ does not contain the witness e_i of P_i or any ancestor of e_i, the algorithm adds e_i to F'. The algorithm performs no further tasks in the iteration for i. Note that $1 \in I$ always holds. If $i > 1$, the algorithm proceeds to the next iteration. If $i = 1$, the deletion phase terminates, and the algorithm outputs F'.

Lemma 6. *Suppose that the algorithm outputs an edge set F'. Then it satisfies the following conditions.*

(i) *$F' \cap E_j \neq \emptyset$ for each $j \in [k]$.*

(ii) *Let $i \in I$. Each subpath between an end node of P_i and lca_i includes at most one edge in F'.*

(iii) *Let $i, i' \in I$ with $i \neq i'$. If an edge in $F' \cap E_{i'}$ is incident to a node v in V_i, and if $\mu(v,i) > 0$, then $F' \cap E_i$ contains an edge incident to v.*

Proof. To prove (i), we first consider the case of $j \in I$. The deletion phase adds an edge $e \in E_j$ to F' before completing the iteration for j. If e is removed from F' in a later iteration for j', then $e \in E_{j'}$ and the algorithm adds an ancestor $f \in E_{j'}$ of e to F'. $\mathrm{lca}_{j'}$ is a descendant of lca_j because $j' < j$ and $e \in E_{j'} \cap E_j$. Hence, f also covers P_j. Even if f is eventually removed from F', we can similarly prove that P_j is covered by another edge that is added to F' in the iteration. Hence, at completion of the algorithm, P_j is covered by F' if $j \in I$.

Before discussing the case of $j \notin I$, we note that if $j \in I$, then $F' \cap E_j$ contains the witness e_j of P_j or one of its ancestors at completion of the algorithm. This is guaranteed by the deletion phase during iteration for j, which adds e_j to F' if F' does not contain e_j or any of its ancestors.

We now discuss the case of $j \notin I$. By definition of the increase phase, F contains an edge e that covers P_j, and e is the witness of P_i for some $i \in I$ with $i < j$. By the above observation, F' contains e or one of its ancestors $f \in E_i$. f also covers P_j, because $i < j$. Hence, F' contains an edge covering P_j even if $j \notin I$.

Next, we prove (ii) by induction on i. Let us consider the case of maximum i in I. The first iteration of the deletion phase adds the witness e_i of P_i to F'. Suppose that another edge $f \in E_i$ is added to F' in the iteration for $i' \in I$ in the deletion phase. If f is added to F before the iteration for i in the increase phase, then i does not belong to I, which is a contradiction. Otherwise, f is added to F because it is incident to an end node of e, and e does not cover $E_{i'}$. Since $i' < i$, this implies that either f is a descendant of e or $u_e = u_f$. In the former case, e is not chosen as the witness of P_i, because the witness is a bottleneck edge of maximum depth with both end nodes

non-relaxable, which contradicts the definition of e. In the latter case, $u_e = u_f = \mathrm{lca}_i$, consistent with (ii) in Lemma 6.

We next consider the case of non-maximal i in I. Let K be the set of edges in the subpath between an end node of P_i and lca_i. Suppose that at the start of the iteration for i in the deletion phase, there exist distinct edges $e, e' \in K \cap F'$. Without loss of generality, we assume that e' is an ancestor of e. Assume that the deletion phase adds e to F' during the iteration for j with $i < j$, and adds e' to F' during the iteration for j' with $i < j'$. Since lca_i is a descendant of lca_j, e' covers P_i; therefore, e' also covers P_j. Hence, the subpath of P_j between l_e and lca_j is covered by two edges in F', which is a contradiction by induction. Therefore, at the beginning of the iteration for i in the deletion phase, K is covered by at most one edge in F'.

At the start of the iteration for i, suppose that K is covered by an edge $e \in F'$. If another edge $f \in K$ is added to F' during this iteration, then f is incident to a node $v \in \tilde{V}_i$, and another edge f' incident to v is in F'. Suppose that f' is added to F' during the iteration for j in the deletion phase. If v is an ancestor of e, then e is removed from F' when f is added to F'. If v is a descendant of e, then P_j is covered by e because $j > i$. Therefore, the subpath of P_j between $l_{f'}$ and lca_j is covered by both e and f', which again is a contradiction by induction.

K could also be covered by two edges if the witness e_i of P_i is in K and is added to F' in the iteration for i, even though one of its descendants has already been in F'. We now demonstrate that this situation does not occur. Note that each edge in $F \cap E_i$ is either e_i or is incident to a node in $V_{i'}$ for some $i' \in I$ with $i' > i$. Since $P_{i'}$ is not covered by e_i and lca_i is not an ancestor of $\mathrm{lca}_{i'}$, e_i is not an ancestor of any edge in $F \cap E_i$. Hence, no descendant of e_i is added to F'.

Finally, we prove (iii). Suppose that some $i, i' \in I$ violates the claim of (iii). We consider such a pair of i and i' that minimizes $|i - i'|$. Let $e \in F' \cap E_{i'}$ be an edge incident to a node $v \in V_i$. Suppose that $i' > i$. If $F' \cap E_i$ contains no edge incident to v, then the iteration for i during the deletion phase adds an ancestor edge of v in E_i to F'. This edge also covers $P_{i'}$ because $\mathrm{lca}_{i'}$ is an ancestor of lca_i, which contradicts (ii). Next, we suppose that $i' < i$. Claim (iii) is obvious when e covers P_i; thus, we suppose that e does not cover P_i. $\mu(v, i) > 0$ indicates that e was not in F at the start of the iteration for i in the increase phase. Let $j \in I$ be the index for which e is added to F during the iteration for j in the increase phase, where $j \geq i$. By the definition of the increase phase, the witness of P_j is incident to v, and because $\mu(v, i) > 0$, an edge incident to v in E_i is added to F simultaneously with e. Since this edge is not in F', the iteration for i in the increase phase adds an ancestor edge of v in E_i to F'. This edge covers P_j, indicating that the subpath of P_j between its one end node and lca_j is covered by two edges if $j > i$, which contradicts (ii). Therefore, $j = i$. However, e is not added to F' unless the witness of P_i is added to F' and a contradiction arises. □

Proof of Theorem 2. Let F' denote the edge set output by the algorithm. By claim (i) of Lemma 6, F' is a multicut. Since $\sum_{i \in [k]} \xi(i)$ is a lower bound on the optimal value, it suffices to show that $\sum_{e \in F'} w(e) + \sum_{v \in V(F')} w(v) \leq 2 \sum_{i \in [k]} \xi(i)$.

Lemma 5 implies that $w(e) = \sum_{i \in [k]: e \in E_i} \nu(e, i)$ for each $e \in F$, and that $w(v) = \sum_{i \in [k]: v \in V_i} \mu(v, i)$ for each $v \in V(F)$. Recall that $F' \subseteq F$. Hence,

$$\sum_{e \in F'} w(e) + \sum_{v \in V(F')} w(v) = \sum_{i \in [k]} \left(\sum_{e \in F' \cap E_i} \nu(e, i) + \sum_{v \in V(F') \cap V_i} \mu(v, i) \right). \tag{9}$$

(iii) of Lemma 6 indicates that $V(F') \cap V_i \subseteq V(F' \cap E_i)$ holds if $i \in I$. If $i \notin I$, then $\nu(e,i) = 0$ for any $e \in E_i$ and $\mu(v,i) = 0$ for any $v \in V_i$. Hence, the right-hand side of (9) is equal to

$$\sum_{i \in I} \sum_{e=uv \in F' \cap E_i} (\nu(e,i) + \mu(v,i) + \mu(u,i)) = \sum_{i \in I} |F' \cap E_i| \xi(i).$$

$|F' \cap E_i| \leq 2$ for each $i \in I$ by (ii) of Lemma 6. Recall that $\xi(i) = 0$ for each $i \notin I$. Therefore, the right-hand-side of (9) is at most $2 \sum_{i \in [k]} \xi(i)$. □

6 Conclusion

In this paper, we emphasized a large integrality gap when the natural LP relaxation is applied to the graph covering problem that minimizes node weights. We then formulated an alternative LP relaxation for graph covering problems in edge- and node-weighted graphs that is stronger than the natural relaxation. This relaxation was incorporated into an exact algorithm for the prize-collecting EDS problem in trees, and a 2-approximation algorithm for the multicut problem in trees. The approximation guarantees for these algorithms match the previously known best results for purely edge-weighted graphs. In many other graph covering problems, the integrality gap in the proposed relaxation would increase if node weights were introduced, because the problems in node-weighted graphs admit stronger hardness results, as shown in the Appendix. Nonetheless, the proposed relaxation is a potentially useful tool for designing heuristics or using IP solvers to solve the above problems.

Acknowledgements

This work was partially supported by Japan Society for the Promotion of Science (JSPS), Grants-in-Aid for Young Scientists (B) 25730008. The author thanks an anonymous referee of SWAT 2014 for pointing out an error on the reduction from the prize-collecting multicut problem to the multicut problem in an earlier version of this paper.

References

[1] M. Bateni, M. Hajiaghayi, and V. Liaghat. Improved approximation algorithms for (budgeted) node-weighted Steiner problems. In *ICALP (1)*, vol. 7965 of *Lecture Notes in Computer Science*, pages 81–92, 2013.

[2] A. Berger, T. Fukunaga, H. Nagamochi, and O. Parekh. Approximability of the capacitated b-edge dominating set problem. *Theoretical Computer Science*, 385(1-3):202–213, 2007.

[3] A. Berger and O. Parekh. Linear time algorithms for generalized edge dominating set problems. *Algorithmica*, 50(2):244–254, 2008.

[4] A. Berger and O. Parekh. Erratum to: Linear time algorithms for generalized edge dominating set problems. *Algorithmica*, 62(1-2):633–634, 2012.

[5] J. Byrka, F. Grandoni, T. Rothvoß, and L. Sanità. Steiner tree approximation via iterative randomized rounding. *Journal of the ACM*, 60(1):6, 2013.

[6] R. D. Carr, T. Fujito, G. Konjevod, and O. Parekh. A $2\frac{1}{10}$-approximation algorithm for a generalization of the weighted edge-dominating set problem. *Journal of Combinatorial Optimization*, 5(3):317–326, 2001.

[7] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating multicut and sparsest-cut. *Computational Complexity*, 15(2):94–114, 2006.

[8] C. Chekuri, A. Ene, and A. Vakilian. Prize-collecting survivable network design in node-weighted graphs. In *APPROX-RANDOM*, vol. 7408 of *Lecture Notes in Computer Science*, pages 98–109, 2012.

[9] V. Chvátal. A greedy heuristic for the set-covering problems. *Mathematics of Operations Research*, 4:233–235, 1979.

[10] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, 1994.

[11] T. Fujito and H. Nagamochi. A 2-approximation algorithm for the minimum weight edge dominating set problem. *Discrete Applied Mathematics*, 118(3):199–207, 2002.

[12] T. Fukunaga. Spider covers for prize-collecting network activation problem. *CoRR*, abs/1310.5422, 2013.

[13] N. Garg, V. V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996.

[14] N. Garg, V. V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, 1997.

[15] M. X. Goemans, N. Olver, T. Rothvoß, and R. Zenklusen. Matroids and integrality gaps for hypergraphic steiner tree relaxations. In *STOC*, pages 1161–1176, 2012.

[16] M. Hajiaghayi, V. Liaghat, and D. Panigrahi. Online node-weighted Steiner forest and extensions via disk paintings. In *FOCS*, pages 558–567, 2013.

[17] D. S. Hochbaum. Heuristics for the fixed cost median problem. *Mathematical Programming*, 22(1):148–162, 1982.

[18] N. Kamiyama. The prize-collecting edge dominating set problem in trees. In P. Hlinený and A. Kucera, editors, *Mathematical Foundations of Computer Science 2010, 35th International Symposium, MFCS 2010, Brno, Czech Republic, August 23-27, 2010. Proceedings*, vol. 6281 of *Lecture Notes in Computer Science*, pages 465–476. Springer, 2010.

[19] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, 2008.

[20] P. N. Klein and R. Ravi. A nearly best-possible approximation algorithm for node-weighted Steiner trees. *Journal of Algorithms*, 19(1):104–115, 1995.

[21] J. Könemann, S. S. Sadeghabad, and L. Sanità. An LMP $O(\log n)$-approximation algorithm for node weighted prize collecting Steiner tree. In *FOCS*, pages 568–577, 2013.

[22] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383–390, 1975.

[23] A. Moss and Y. Rabani. Approximation algorithms for constrained node weighted Steiner tree problems. *SIAM Journal on Computing*, 37(2):460–481, 2007.

[24] J. Naor, D. Panigrahi, and M. Singh. Online node-weighted Steiner tree and related problems. In *FOCS*, pages 210–219, 2011.

[25] Z. Nutov. Approximating Steiner networks with node-weights. *SIAM Journal on Computing*, 39(7):3001–3022, 2010.

[26] Z. Nutov. Approximating minimum-cost connectivity problems via uncrossable bifamilies. *ACM Transactions on Algorithms*, 9(1):1, 2012.

[27] Z. Nutov. Survivable network activation problems. In *LATIN*, vol. 7256 of *Lecture Notes in Computer Science*, pages 594–605, 2012.

[28] D. Panigrahi. Survivable network design problems in wireless networks. In *SODA*, pages 1014–1027, 2011.

[29] O. Parekh. Approximation algorithms for partially covering with edges. *Theoretical Computer Science*, 400(1-3):159–168, 2008.

[30] A. Vakilian. Node-weighted prize-collecting survivable network design problems. Master's thesis, University of Illinois at Urbana-Champaign, 2013.

[31] V. V. Vazirani. *Approximation algorithms*. Springer, 2001.

A EDS problem in general graphs

A.1 Hardness results

In this section, we discuss the EDS problem in general graphs. First, we show that the EDS problem in node-weighted graphs is as hard as the set cover problem, and the EDS problem in edge- and node-weighted graphs is as hard as the (non-metric) facility location problem. To this end, we reduce the set cover problem or the facility location problem to the EDS problem in bipartite graphs. Accordingly, our hardness results hold for the EDS problem in bipartite graphs.

We now define the set cover problem and facility location problem. In the set cover problem, we are given a set V, family $\mathcal{S}$ of subsets of V, and cost function $c\colon \mathcal{S} \to \mathbb{R}_+$. A solution is a subfamily $\mathcal{X}$ of $\mathcal{S}$ such that $\cup_{X\in\mathcal{X}} X = V$. The objective is to minimize $\sum_{X\in\mathcal{X}} c(X)$. Inputs in the facility location problem are a client set V, a facility set F, opening costs $o\colon F \to \mathbb{R}_+$, and connection costs

$d\colon V \times F \to \mathbb{R}_+$. A solution to this problem is a pair of $F' \subseteq F$ and $\rho\colon V \to F'$, and the objective is to minimize $\sum_{f \in F'} o(f) + \sum_{v \in V} d(v, \rho(v))$.

Given an instance of the set cover problem, define a facility set F whose members each corresponds to a set in $\mathcal{S}$; We let X_f denote the member of $\mathcal{S}$ corresponding to $f \in F$. Define opening costs o and connection costs d such that $o(f) = c(X_f)$, $d(v, f) = 0$ if $v \in X_f$, and $d(v, f) = +\infty$ otherwise. Then, the instance (V, F, o, d) of the facility location problem is equivalent to the instance $(V, \mathcal{S}, c)$ of the set cover problem, implying that the non-metric facility location problem generalizes the set cover problem.

Theorem 4. *If the EDS problem in node-weighted bipartite graphs admits a γ-approximation algorithm, then the set cover problem also admits a γ-approximation algorithm. If the EDS problem in edge- and node-weighted bipartite graphs admits a γ'-approximation algorithm, then the facility location problem also admits a γ'-approximation algorithm.*

Proof. To prove Theorem 4, we reduce the facility location problem to the EDS problem. From a client set V and facility set F, we construct the complete bipartite graph with bipartition V and F. Moreover, for each $v \in V$, we add a new node v' and an edge e_v that joins v and v'. Note that this operation retains the bipartite property of the graph. We define edge weights w such that $w(e) = d(v, f)$ for each $e = vf$, where $v \in V$ and $f \in F$ and $w(e_v) = 0$ for each $v \in V$. The node weights w are defined by $w(v) = 0$ and $w(v') = +\infty$ for each $v \in V$, and $w(f) = c(f)$ for each $f \in F$.

Let S be an EDS for this graph. We can assume $e_v \notin S$ because $w(v') = +\infty$. Since S dominates each e_v, S contains an edge incident to each $v \in V$. Moreover, if S contains more than one edge incident to $v \in V$, S remains an EDS if one of these edges is arbitrarily discarded. Hence S contains exactly one edge incident to $v \in V$, which joins v and a node $f \in F$. Let $\rho(v)$ be the opposite end node of the edge in S incident to v, and let $F' = \{\rho(v)\colon v \in V\}$. Then, (F', ρ) is a solution to the facility location problem, with cost equaling the weight of S. Conversely, given a solution to the facility location problem, define S as $\{v\rho(v)\colon v \in V\}$. Then S is an EDS of the graph, and the weight of S equals the cost of the solution to the facility location problem. Hence, the reduction is that required in the latter part of the theorem.

As observed above, the set cover problem corresponds to the instances of the facility location problem with zero connection cost. For these instances, the reduction defines instances of the EDS problem in which all edges are weighted zero. The former part of the theorem follows from this statement. □

Without describing the details, we can also show that reductions in the proof of Theorem 4 are applicable (with modification) to fundamental covering problems such as the Steiner tree problem, T-join problem, and edge cover problem.

A.2 $O(\log |V|)$-approximation algorithm

Since it is NP-hard to achieve the $o(\log |V|)$-approximation in the set cover problem, by Theorem 4, the same situation exists for the EDS problem in node-weighted bipartite graphs. Here, we propose an $O(\log |V|)$-approximation algorithm for the prize-collecting EDS problem in general edge- and node-weighted graphs.

Our algorithm reduces the prize-collecting EDS problem to the edge cover problem. Given an undirected graph $G = (V, E)$ and a set U of nodes, we define an edge cover as a set F of edges such that $\delta_F(v) \neq \emptyset$ for each $v \in U$. Given $w\colon E \cup V \to \mathbb{R}_+$, the problem seeks an edge cover F that minimizes $w(F) + w(V(F))$. For an instance $I' = (G, U, w)$ of the edge cover problem, we apply the following LP relaxation:

$$
\begin{array}{lll}
EC(I') = & & \\
\text{minimize} & \displaystyle\sum_{e\in E} w(e)x(e) + \sum_{v\in V} w(v)x(v) & \\
\text{subject to} & \displaystyle\sum_{e\in\delta(v)} x(e) \geq 1 & \text{for } v \in U, \\
& x(v) \geq x(e) & \text{for } v \in V, e \in \delta(v), \\
& x(e) \geq 0 & \text{for } e \in E, \\
& x(v) \geq 0 & \text{for } v \in V.
\end{array}
$$

Lemma 7. *Suppose that an algorithm computes an edge cover of weight at most $\gamma EC(I')$ for any instance I' of the edge cover problem. Then the prize-collecting EDS problem in edge- and node-weighted graphs admits a 4γ-approximation algorithm.*

Proof. We first solve $P(I)$ for the given instance $I = (G, w, \pi)$ of the prize-collecting EDS problem. Let (x, y, z) be the obtained optimal solution to $P(I)$. Recall that $\mathcal{E} = \{\delta(e)\colon e \in E\}$ in the EDS problem; thus we denote $y(C, e)$ and $z(C)$ by $y(e', e)$ and $z(e')$ respectively, where e' is the edge corresponding to $C \in \mathcal{E}$. Let $U = \{v \in V\colon \sum_{e\in\delta(v)} x(e) \geq 1/4\}$. An instance I' of the edge cover problem is assumed to consist of G, U, and w.

We now prove that $EC(I') \leq 4(\sum_{e\in E} w(e)x(e) + \sum_{v\in V} w(v)x(v))$. In the following, $4x$ is denoted by x'. Since $\sum_{e\in\delta(v)} x(e) \geq 1/4$ for each $v \in U$, $\sum_{e\in\delta(v)} x'(e) \geq 1$ for each $v \in U$, implying that x' satisfies the first constraints of $EC(I')$. By the minimality of x, there exists $e' \in \delta(e)$ such that $x(e) = y(e', e)$ for any $e \in E$. The second constraint for v and e' in $P(I)$ implies that $x'(v) = 4x(v) \geq 4y(e', e) = 4x(e) = x'(e)$ for $e \in E$ and an end node v of e. Hence, x' satisfies the second constraints of $EC(I')$. Therefore, x' is a feasible solution to $EC(I')$, and $EC(I')$ is at most $4(\sum_{e\in E} w(e)x(e) + \sum_{v\in V} w(v)x(v))$.

Let F be an edge cover of weight at most $\gamma EC(I')$ for I'. Let $D = \{e \in E\colon z(e) \leq 1/2\}$. By the first constraint of $P(I)$, $\sum_{e'\in\delta(e)} y(e, e') \geq 1/2$ for each $e \in D$. Let u and v be the end nodes of $e \in D$. Since $\sum_{e'\in\delta(e)} y(e, e') \leq \sum_{e'\in\delta(u)} y(e, e') + \sum_{e'\in\delta(v)} y(e, e') \leq \sum_{e'\in\delta(u)} x(e') + \sum_{e'\in\delta(v)} x(e')$, at least one of u and v is included in U. Note that each node in U has some incident edge in F. Hence, each edge in D is dominated by F. Therefore, F must pay at most $\sum_{e\notin D} \pi(e) \leq 2\sum_{e\in E} \pi(e)z(e)$ as the penalty. Summing up, the objective value of F for I does not exceed

$$\gamma EC(I') + 2\sum_{e\in E} \pi(e)z(e) \leq 4\gamma\left(\sum_{e\in E} w(e)x(e) + \sum_{v\in V} w(v)x(v)\right) + 2\sum_{e\in E} \pi(e)z(e) \leq 4\gamma P(I).$$

Since $P(I)$ is a lower bound on the optimal value of I, F achieves a 4γ-approximation for I. □

To obtain an $O(\log |V|)$-approximation algorithm for the prize-collecting EDS problem, it suffices to obtain an algorithm for the edge cover problem required in Lemma 7 with $\gamma = O(\log |V|)$. As a sequel, we observe that such an algorithm indeed exists.

Since each node in U is included in $V(H)$ for any edge cover H, an algorithm that solves instances of $w(v) = 0$ for all $v \in U$ is sufficient. Moreover, if G contains an edge e joining nodes u and v in U, the instance is transformed into an equivalent instance by inserting a new node s that subdivides the edge uv and setting the new node and edge weights as $w(us) = w(sv) = 0$ and $w(s) = w(e)$. In the obtained instance, U forms an independent set. Such instances of the edge cover problem are included in the (non-metric) facility location problem. In fact, we may regard U and $V \setminus U$ as the client and facility sets, respectively. The weights of edges between clients and facilities represent connection costs, and the weights of clients indicate opening costs. Each edge in edge covers joins a client and facility and naturally allocates the client to the facility.

Consider an instance of the facility location problem. We define a *star* as a set comprising a facility $f \in F$, and clients $v_1, \ldots, v_k \in V$. The cost of the star is $o(f) + \sum_{i=1}^{k} d(v_i, f)$. Let $\mathcal{S}$ be the set of all stars. Identifying the star as the subset $\{v_1, \ldots, v_k\}$ of V, the facility location problem can be regarded as the set cover problem with the set V and family $\mathcal{S}$ of subsets of V. Hence, the greedy algorithm for the set cover problem achieves $O(\log |V|)$-approximation for the facility location problem. An analysis of this greedy algorithm [22, 9] has shown that the costs of solutions are bounded by an LP relaxation of the set cover problem (see also [31]). Our LP relaxation $EC(I')$ applied to an instance I' of the edge cover problem is equivalent to this LP relaxation for the set cover problem derived from I'. The runtime of the greedy algorithm is a primary concern, because the size of $\mathcal{S}$ in the set cover problem is not bounded by a polynomial of the input size of I'. The greedy algorithm can operate in polynomial time for the given instances in the facility location problem and edge cover problem [17]. Therefore, we can state the following theorem.

Theorem 5. *The prize-collecting EDS problem in general edge- and node-weighted graphs admits an $O(\log |V|)$-approximation algorithm.*